

215-TP-001-001

**NASA's Earth Observing System Data and Information
System (EOSDIS): An Integrated System for Processing,
Archiving, and Disseminating High-Volume Earth Science
Imagery and Associated Products**

**Technical Paper--Not intended for formal review or
government approval.**

July 1996

Prepared Under Contract NAS5-60000

RESPONSIBLE ENGINEER

<u>Stephen Fox, Narayan Prasad, Marti Szczur /s/</u>	<u>10/11/96</u>
Stephen Fox, Narayan Prasad, Marti Szczur (NASA)	Date
EOSDIS Core System Project	

SUBMITTED BY

<u>Robert Clinard /s/</u>	<u>10/11/96</u>
Robert Clinard	Date
EOSDIS Core System Project	

Hughes Information Technology Systems
Upper Marlboro, Maryland

This page intentionally left blank.

NASA's EARTH OBSERVING SYSTEM DATA AND INFORMATION SYSTEM (EOSDIS): AN INTEGRATED SYSTEM FOR PROCESSING, ARCHIVING, AND DISSEMINATING HIGH-VOLUME EARTH SCIENCE IMAGERY AND ASSOCIATED PRODUCTS

Stephen Fox, Narayan Prasad
Hughes Information Technology Systems
1616 McCormick Drive
Upper Marlboro, MD 20774
(sfox@eos.hitc.com, nprasad@eos.hitc.com)

Marti Szczur
NASA/Goddard Space Flight Center
Greenbelt, MD 20771
(marti.szczur@gsfc.nasa.gov)

Abstract

The Earth Observing System Data and Information system (EOSDIS) is being developed by NASA's Goddard Space Flight Center to acquire a comprehensive, global, 15-year data set containing several petabytes of data. EOSDIS will command and control the series of EOS satellites, as well as provide distributed data processing, archival and distribution services across eight service provider sites. This paper provides an overview of the core components of EOSDIS focusing on the system capabilities provided for planning and processing science data products.

Earth Observing System Data and Information System (EOSDIS)

Barely a decade ago, concerns such as Ozone depletion, deforestation and global warming interested only a handful of atmospheric and environmental scientists. Today, increasing evidence of large-scale change has made us aware of the threats to our life-sustaining global environment. Within the span of a few human generations, we have changed our world significantly without understanding the long-term effects on its ability to sustain life. Motivated by this critical concern, the U.S. Government initiated the Global Change Research Program to develop a predictive understanding of the global environment. Scientific success in understanding global environmental change depends on integration and management of numerous

data sources, extensive data holdings and heterogeneous data products.

In response to this challenge and with the unique contribution of providing a global view from space, NASA initiated the first comprehensive study of interaction between atmosphere, oceans, land and life on Earth, called the Earth Observing System (EOS)^{1,2}. EOS is a long term program that will study the Earth using interrelated data from advanced instruments on-board a series of Earth-orbiting satellites. Each EOS spacecraft will focus on a different aspect of global climate change and complex interactions within the Earth's environment. The first satellite is to be launched in 1998, and will observe the Earth's surface, clouds, aerosols and radiation balance. The EOS schedule targets a minimum of 15 years of continuous Earth observations.

NASA's Goddard Space Flight Center is developing the EOS Data and Information System (EOSDIS) to manage the data

resulting from the EOS satellites, field measurement programs, and other data essential for the interpretation of these measurements. Besides providing command and control of the EOS spacecraft, it will process, store and distribute the data to thousands of international scientists and other users. The EOSDIS will serve as the connecting link between the observations made from space and the scientists who will use these data in their research.

NASA historically has built separate and centralized data and information systems for each science mission because of past focus on discipline-oriented research and the difficulties in building open, distributed systems. Modern technology now allows broad coupling of expanded services with advanced mission data collection to support interdisciplinary Earth science investigations, a necessary condition for resolving issues associated with global change.

EOSDIS' infrastructure, the EOSDIS Core System (ECS), provides EOS and other U.S. and international scientists a broad range of services from 8 science data centers, the Distributed Active Archive Centers (DAACs), operated by NASA and other agencies. The ECS infrastructure also supports exchange of data and research results within the science community, across multiple agencies and internationally, through the broader Global Change research community (GCDIS) as well as educational and commercial users (UserDIS). ECS is the evolutionary base for accelerating the pace of Earth science research.

EOSDIS Ground System Architecture

The ECS architecture^{4,5,9} provides the base for an open-system design with a low cost of changing services and data types. Services and new data products can be added incrementally to the system over time, providing low cost of entry for all service providers. The service providers include Value-Added Providers (VAPs), which

provide unique services beyond those in ECS, science computing facilities (SCFs), which provide the core Earth science data processing algorithms and conduct ongoing research, and all other related data centers.

The subsystems illustrated in Figure 1 capture the complete distributed functionality of the ECS spread across the eight sites illustrated in the map. The subsystems are:

CLIENT - Provides the "client" part of the "Client / Server" access paradigm through graphical user interface and data/service access tools, as well as application program interface (API) libraries to ECS services.

INTEROPERABILITY - Provides application level "middle-ware" which facilitates dynamic client access to services and providers holding data collections and services.

DATA MANAGEMENT - Provides system wide distributed search and access services with multiple science discipline views of data collections and "one stop shopping" with location transparent access to those services and data.

DATA SERVER - Provides locally optimized search, access, archive and distribution services with a science discipline view of data collections and an extensible Earth Science Data Type and Computer Science Data Type view of the archive holdings.

INGEST - Provides the "clients" for the importation of data (science products, ancillary, correlative, documents, etc.) into ECS data repositories (Data Servers) on an ad hoc or scheduled basis and deals with external system interfaces.

PLANNING - Provides for pre-planning of routine/ad hoc/on-demand science data processing as well as management

functions for handling deviations from the operations plans for individual DAAC sites.

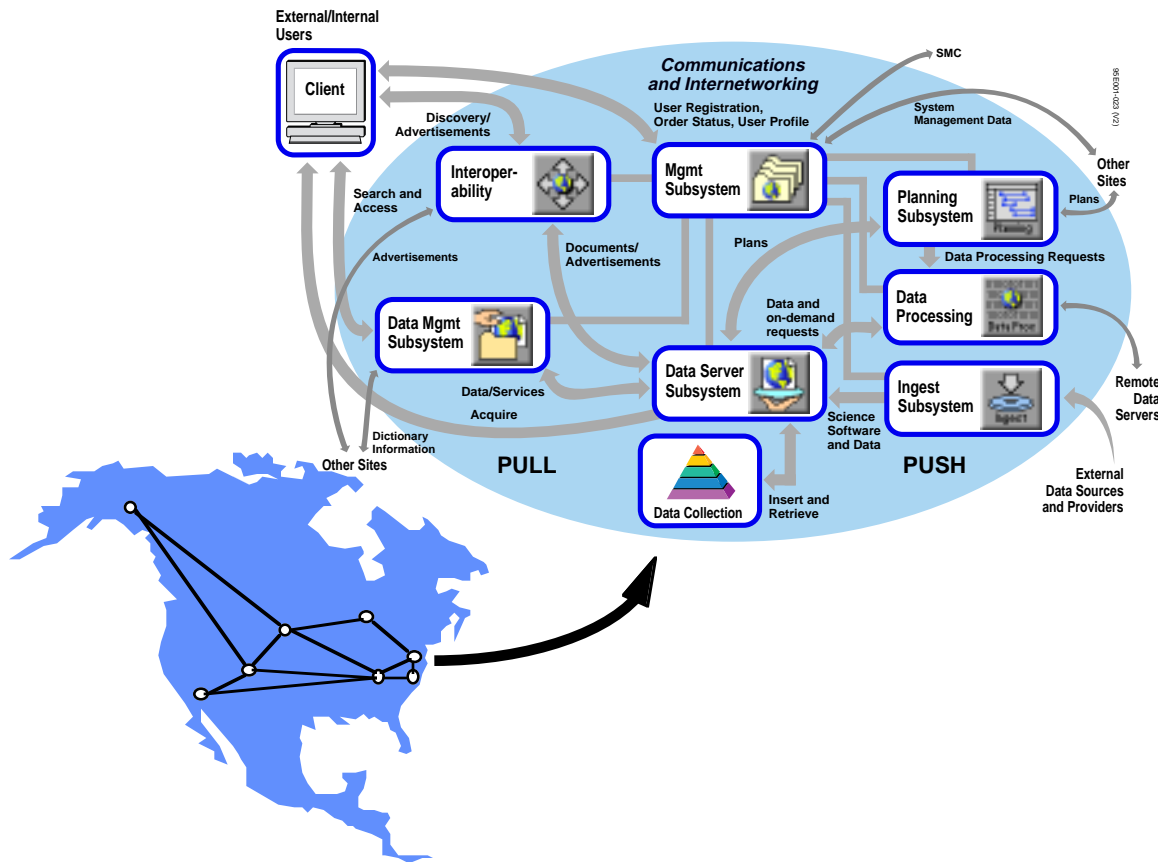


Figure 1 ECS Subsystem Level Architecture

DATA PROCESSING - Provides the functions to host science algorithm software, perform science software integration and test, data processing, process resource management and includes facilities and toolkits which offer true software portability across advanced computing platforms.

MANAGEMENT - Provides functions for system startup/shutdown, resource management, performance monitoring, error logging, system and science software configuration management, and resource accounting.

COMMUNICATIONS - Provides the distributed computing infrastructure that enables intra- and inter-site communications between the subsystems.

From a scientist user's perspective, the ECS infrastructure appears as services, and interfaces to services, which are displayed on the user's workstation "desktop". This perspective is similar to the view of the World Wide Web (WWW) servers as seen through a local client such as Mosaic or Netscape. The Client Subsystem allows users to easily explore and locate services

that are "advertised" by data producers through the Interoperability Layer. Advertisements via the Interoperability Subsystem provide integrated views of services, data sets, and providers. After identifying a service of interest, users can immediately invoke the provider service, or the reference to the provider service can be saved on the desktop for later use. For example, services are available that support cross-DAAC searching, as well as, searching of local data collections at individual DAACs. The determination of which system services to use to satisfy a user service request is made automatically by EOSDIS and is transparent to the user. Service providers are not limited to the 8 DAACs but can access another data center, scientist, or value added service provider who adopts ECS interfaces and protocols and can advertise their services through the Interoperability Layer.

ECS's open architecture allows provider autonomy and independent, evolutionary development of components to improve services offered to users. Because of logical distribution, users and providers can tailor their components to their environments and still operate together as an extended distributed system such as found with the World Wide Web model. Evolutionary upgrades will be implemented depending on technology and cost considerations. For example, content-based searching of EOS products is not practical with current technology, but is a candidate for a future evolutionary upgrade.

Science Data Processing Challenges

ECS will produce 260 standard data products on an ongoing basis. These standard data products typically consist of geophysical parameters derived from calibrated EOS instrument data and other previously derived data. These products are produced from raw instrument data which is received from the EOSDIS ground station.

On a daily basis, ECS must process over 480 gigabytes of raw data and produce 18,000 product instances with an aggregate size of 1,600 gigabytes. Thus, ECS data archives will grow at the rate of two terabytes of data per day and are expected to reach a petabyte in size by 1999.

ECS provides system capabilities for integrating, planning, and executing science software that will produce these standard data products. Each EOS instrument has an associated Instrument Team of scientists and EOS investigators who develop the science software that produces the standard data products. Instrument Teams and EOS investigators are located at over sixty Science Computing Facilities (SCFs). After development at an SCF, the science software is then delivered to a DAAC where it is tested and then integrated into the ECS production environment. The distributed, multi-mission nature of EOS presents a number of challenges in the science data processing area:

- ECS must be able to efficiently integrate and test science software developed on different computing platforms at different SCFs.
- As investigator's improve their scientific understanding, they will want to refine their science software and reprocess a potentially large portion of ECS data holdings using updated algorithms.
- Because investigators want to take advantage of the complementary nature of the EOS instrument suite, there is a high degree of data dependence among the standard data products. Also, with standard data production taking place at six different locations, many data dependencies involve cross-site interactions. Thus, ECS must provide a toolset capable of supporting cross-DAAC planning and scheduling.
- In addition to standard production, ECS supports on-demand production of

certain products in response to user requests. The planning and scheduling functions must be able to interleave on-demand production requests with standard production in a manner that maximizes computing resources and maintains the overall production schedule.

The remaining subsections describe the ECS Science Software Integration and Test, Planning, and Data Processing capabilities and discuss how they address these challenges.

Science Software Integration and Test

SCFs located at EOS Investigator home institutions are responsible for developing, validating and maintaining the science algorithms and software that produce the standard data products. Over the mission lifetime, each SCF will make a number of science software deliveries to ECS as science algorithms are initially created and subsequently refined. Thus, a key requirement of ECS is to integrate new or updated science software into the ECS production environment in an efficient and reliable manner. In order to accomplish this, ECS provides two key components: a Science Data Processing Toolkit which encapsulates the interface to the ECS system and a suite of science software integration and test tools which support configuration management and testing of the science software at the DAAC.

Science Data Processing (SDP) Toolkit The SDP Toolkit⁷ consists of a set of fully tested and reliable C and FORTRAN language functions, customized and optimized for application to ECS. The SDP Toolkit provides the interface between the science software and the ECS system, including its processing environment, system

management, and data management functions. It promotes the ability to port the science software to different platforms at the DAAC. It also reduces duplication of coding effort among the disparate science software developers by providing a set of scientific functions commonly used by all science software.

Two versions of the SDP Toolkit are available. The SCF version is designed to operate in a standalone mode on a variety of computing platforms. It provides a set of interfaces that encapsulate operating system, file handling, error reporting, metadata generation, and common scientific functions. SCFs develop and test their science software against these interfaces at their local facilities. The DAAC version of the SDP Toolkit provides interfaces identical to the SCF version, however in the DAAC version, the interfaces are implemented by ECS system functions. This approach greatly improves science software portability as most platform-specific software is encapsulated within the SDP Toolkit, thus insulating the science software from platform-dependent system software.

The SDP Toolkit provides two types of tools: mandatory tools, which are required to be used in science software, and are checked during DAAC integration and test; and optional tools, which are common utilities provided with the intention of saving individual SCF development effort, as well as reducing duplication of effort across the different SCFs. These tools are summarized in Table 1.

Science Software Integration and Test (SSI&T) Tools SSI&T tools⁸ are used at the DAAC in order to facilitate the transition of the science software in the ECS production environment. The tools

Table 1 Tools Provided in SDP Toolkit

Mandatory Tools	Description
Generic Input/Output Tools	Provide the means to open and close temporary and intermediate duration files
Error and Status Message Tools	Provide general error reporting and status log messaging capabilities interfacing to ECS communications and management services.
Process Control Tools	Provide the primary interface to the ECS Data Processing subsystem. The major use of these tools is to access physical filenames and file attributes. In addition, they retrieve science software defined parameters used in processing (e.g., selecting between variations of an algorithm that have been included in the source code)
HDF Access Tools	Read and write standard data product files in ECS standard format.
Level 0 Access Tools	Read raw Level 0 data from EOS spacecraft
Metadata Access Tools	Read, alter, write, and append metadata to standard data product files.
EOS Spacecraft Ephemeris and Attitude Tools	Read ephemeris and attitude data
Time and Date Tools	Perform time and date conversions between selected time systems
Optional Tools	Description
Ancillary Data Access Tools	Provide access to ancillary data such as NOAA NMC data and Digital Elevation Model data
Celestial Body Position Tools	Locate the spacecraft, sun and other celestial bodies.
Coordinate System Conversion Tools	Allow celestial reference frame coordinate conversions and perform related tasks such as locating the sub-satellite point.
Constant and Unit Conversion Tools	Allow access to physical constants and unit conversions.
IMSL Tools	Provide mathematical and statistical functions.

are mostly Commercial Off-The Shelf (COTS) products and are summarized in Table 2.

The DAAC Integration and Test (I&T) team uses these tools to implement a well defined process for integrating and testing science software before placing it into the production environment. In phase one of this process, the I&T team receives science software deliveries from an SCF either electronically or via media. After receipt, a science software delivery is examined for correctness and completeness. This includes examining accompanying documentation,

verifying that prescribed coding standards have been followed, and running preliminary static and dynamic diagnostic tools to check for potential errors in the source code. Once a delivery is found to be correct and complete, all files are placed under configuration control.

In the second phase of the process, the science software is compiled and linked with the DAAC version of the SDP Toolkit. A series of acceptance tests, supplied by the SCF, are run to verify that the science software executes correctly in the DAAC environment.

If

Table 2 Science Software Integration and Test Tools

Capability	Tool
Receive science software delivery	ftp and hard media ingest
Configuration manage delivered science software	ClearCase
SSI&T problem tracking	Remedy Trouble Ticketing and DDTS
Standards compliance checking	FORcheck (for FORTRAN 77) Lint (for C) Custom developed scripts
View documentation	Softwindows, MS Office Ghostscript/Ghostview (for postscript) NCSA Mosaic (for HTML) Adobe Acrobat (for PDF)
Compile and link delivered source code	SGI, NAG, and KAI compilers and debuggers, SGI CASEVision
Run test cases	Custom developed manager
Examine test outputs, including metadata	Custom developed file comparison tools IDL Custom developed HDF file cracker (EOSView)
Collect resource profiles	CASEVision
Update system databases	Custom developed scripts
Write reports and maintain logs	Softwindows, MS Office

any anomalies are discovered, they are reported to the SCF for correction. If needed, the SCF team can remotely log on to the DAAC integration and test environment to aid in diagnosing a problem. Science software is packaged into one or several Product Generation Executives (PGEs). Each PGE consists of one or more compiled binary executables and/or UNIX shell scripts, and is the smallest unit that is planned and scheduled.

Once the science software is executing correctly, resource utilization statistics are collected for each PGE and stored in the planning database. These utilization statistics capture elapsed execution time, CPU time, memory, and disk space requirements and are used by the Planning Subsystem in projecting the data processing resources required to execute the science software.

In order to reduce the risk of science software integration and test problems, a

pre-operational release of the science software integration and test and science processing environment was delivered to three DAACs in January 1996. This release has been used to gain early experience with science software integration and test utilizing beta algorithms from three instrument teams. As no major installation or integration problems were uncovered with this early release, the concept of using a SDP Toolkit to improve science software portability was validated.

Planning

The Planning subsystem⁶ provides the DAAC with the ability to create, modify, and implement a production plan. As shown in Figure 2, the production plan is generated by expanding production requests (i.e., instructions which specify which products should be produced for what periods of time) into individual data processing requests, using PGE resource profiles, predicted resource availability, predicted

input data availability, and a processing strategy which defines priorities. Multiple candidate plans can be created for consideration, but only one plan is activated.

Planning supports the operations staff in developing a production plan based on a locally defined strategy, reserving the resources to permit the plan to be achieved, and implementing the plan as data and processing requests are received. It also allows the site operations staff to negotiate on a common basis with other DAACs and EOSDIS System management if any change to their production plan causes conflict with other provider sites plans (e.g., where dependencies between processing algorithms cannot be fulfilled). This subsystem provides the functions needed to pre-plan routine data processing, schedule on-demand processing, and dispatch and manage processing requests. The subsystem provides access to the data production schedules at each site, and provides management functions for handling deviations from the schedule to operations and science users.

Production requests are entered by DAAC operations personnel to handle standard processing and reprocessing. These requests are maintained in a planning database and are included in the next candidate plan when it is generated. On-demand requests are accepted by Planning from the Data Server subsystem. On-demand requests are either submitted immediately for processing or added to the planning database for inclusion in the next candidate plan, depending on whether certain predetermined criteria

related to request priority and resource requirements are satisfied.

Some reprocessing requests may involve executing updated science software against a large amount of archived data. In order to ensure that such reprocessing requests do not swamp the system and degrade ongoing production, planning decomposes reprocessing requests into manageable pieces and interleaves them with ongoing production work in a manner that ensures production schedules can be met. If processing requests are delayed because of late arriving dependent data or the failure of a production algorithm, Planning will make use of the temporary availability of computing resources by automatically increasing the number of reprocessing requests to be scheduled for execution. When the delayed processing requests are ready to resume, planning will automatically decrease the number of reprocessing requests in order to allow the production work to catch up.

Data processing requests are automatically submitted by Planning for processing when the required input data are available. For standard processing, input data may not be available when a plan is activated, but is expected to arrive from external data sources during the plan's time frame. Planning is informed of the data availability via data availability notices sent from the Data Server.

Planning tracks the status of all production requests entered and all data processing requests generated. Management reports are generated (either periodically or upon request) with

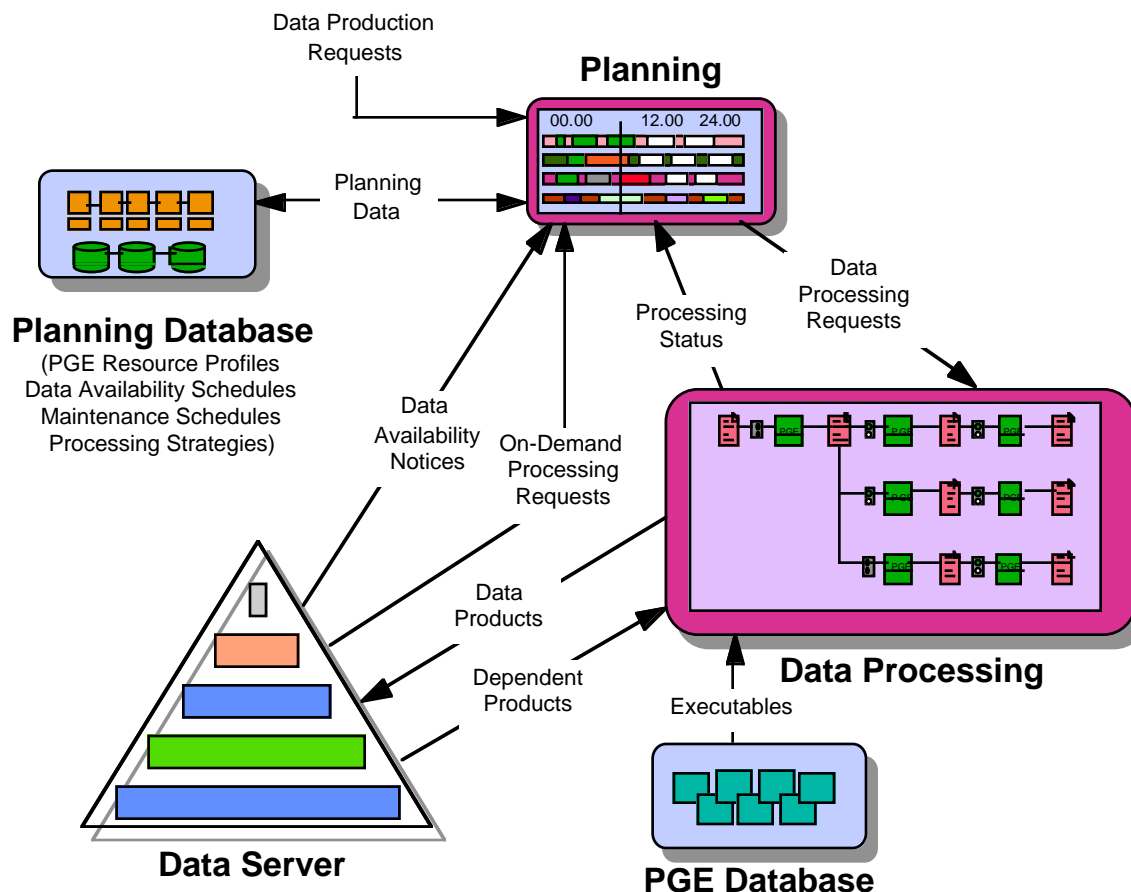


Figure 2 Planning and Data Processing Subsystems

information concerning the planning workload and the status of requests processed.

The Planning subsystem software and hardware must be capable of handling varying workloads across EOSDIS sites. A typical 30 day planning horizon requires analyzing several thousand PGE executions at a small site to over 120,000 PGE executions at the largest site. The Planning subsystem is implemented as a custom developed Sybase SQLServer relational database management system application executing on several Sun workstations and a Sun Ultraserver with multiple CPUs that are sized for the workload at each EOSDIS site.

Data Processing

The Data Processing subsystem³ provides the necessary hardware resources, as well as software for queuing, dispatching and managing the execution of the science software in an environment which is highly distributed and may consist of heterogeneous computing platforms. The subsystem also interacts with the Data Server subsystem to cause the staging and de-staging of data resources in synchronization with processing requirements.

The Data Processing subsystem hardware environment consists of high-end UNIX

servers that execute the science software, RAID-5 disk systems that support the buffering of dependent and newly produced data products, and a scheduling server that supports job queuing, dispatching, and monitoring. Hardware configurations will vary by site with the largest site containing 5 SGI Power Challenge XL servers (total of 48 processors) and over two terabytes of RAID-5 disk storage. The scheduling server at this site must handle a workload of 32,000 jobs per day.

Data processing requests are submitted from the Planning subsystem based on the arrival of dependent data, user requests for on-demand processing, or reprocessing needs. A data processing request contains the information, such as input data identification, output data identification, priority, and hardware resources, that is needed to schedule the execution of a PGE. Generally, a processing request is related to the generation of data products, but these requests may include other types of processing, such as pre-processing of input data, quality assurance processing of generated data products, and possibly resource maintenance.

The Data Processing subsystem uses AUTOSYS, a commercial-off-the-shelf scheduling system from Platinum Technologies. This scheduler is capable of managing the execution of all processing requests in a fully automated fashion according to a pre-planned schedule. The scheduler allocates processing resources to requests according to priority and within the scheduling parameters set by the active plan. It initiates staging and de-staging of data, and re-allocates resources, as necessary, to balance processing load.

During PGE execution, the scheduling system monitors the execution of the PGE and informs the operations staff and the

Planning subsystem of current status. Status may include current processing event history such as data staging activity and the number of job steps completed. If job step failures occur, operations staff are notified and automated error recovery actions are initiated if provided as part of the science software delivery. Error information is automatically captured for a failed PGE and stored in the Data Server subsystem where it can be accessed by science software developers. Typical kinds of error information captured include core dumps, input files, and any temporary files that were produced by the PGE prior to its failure. Upon completion of the execution of a PGE, the scheduling system informs the Planning subsystem and initiates the transfer of any generated data product to the Data Server subsystem.

Summary and Conclusions

The ECS has been designed to provide a comprehensive set of capabilities for managing science data processing in a distributed environment. It provides the capabilities that are necessary to address the key challenges of science data processing in a multi-mission environment.

Integration of science algorithm software, developed by different science software teams on different computing platforms, is facilitated by a Science Data Processing Toolkit that insulates the science software from the underlying computing platform. A comprehensive toolset is provided to aid in science software integration and test prior to inserting the science software into the production environment. Optimization of production resources is provided by a planning system that enables operations personnel to effectively plan routine production of data products, along with on-demand and reprocessing activities. The

system provides the data and tools necessary to enable sites to coordinate cross-site production dependencies and, at the macro level, will dynamically adjust routine production and reprocessing loads to accommodate late arriving dependent data or science software failures. Finally, a scaleable science data processing environment is provided that supports the automated dispatching and monitoring of science software job execution across a distributed set of computing platforms.

The ECS development effort is being managed as multiple incremental releases, each adding new functionality and/or data to the system. Critical Design Reviews have been completed for the first two releases. The first release is scheduled for delivery at the end of 1996 and will support the Tropical Rainfall Measuring Mission (TRMM). The second release is scheduled for delivery in the fall of 1997 and will support the first EOS platform (AM-1) as well as related missions. Additional information about EOSDIS and the ECS design can be found at <http://ecsinfo.hitc.com>.

References

¹Asrar, G. and Dokken, D., "1993 EOS Reference Handbook," August 1993, NASA Headquarters. (Available at http://spso.gsfc.nasa.gov/eos_reference/TOC.html)

²Asrar, G. and Dozier, J., "Science Strategy for the EOS," July 1994, NASA Headquarters.

³ECS Project, "Release B Data Processing Subsystem Design Specification for the ECS Project," March 1996, Hughes Information Technology Systems. (Available at <http://edhs1.gsfc.nasa.gov>)

⁴ECS Project, "Summary of the ECS System Design Specification," June 1994, Hughes Information Technology Systems. (Available at <http://edhs1.gsfc.nasa.gov>)

⁵ECS Project, "Release B Design Specification Overview," March 1996, Hughes Information

Technology Systems. (Available at <http://edhs1.gsfc.nasa.gov>)

⁶ECS Project, "Release B Planning Subsystem Design Specification for the ECS Project," March, 1996, Hughes Information Technology Systems. (Available at <http://edhs1.gsfc.nasa.gov>)

⁷ECS Project, "SDP Toolkit Primer for the ECS Project," April 1995, Hughes Information Technology Systems. (Available at <http://edhs1.gsfc.nasa.gov>)

⁸ECS Project, "Software Developer's Guide to Preparation, Delivery, Integration, and Test with ECS," August 1995, Hughes Information Technology Systems. (Available at <http://edhs1.gsfc.nasa.gov>)

⁹Moxon, B., "Science-based System Architecture Drivers for the ECS Project," Hughes Information Technology Systems. (Available at <http://edhs1.gsfc.nasa.gov>)